SAP with Oracle Real Application Clusters 12c on Microsoft Windows 2012 Release 2

Table of Contents

.

Introduction

Oracle Real Application Clusters 12c (Oracle RAC) and Oracle Automatic Storage Management 12c (Oracle ASM) technologies provide unexcelled availability, scalability and performance for your SAP applications.

This document guides you through the most important steps required, to install such a cluster and will also discuss the recommended path to migrate an Oracle Database 11g Release 2 based system (single instance and RAC) to Oracle RAC 12.1.

The pre-installation steps for setting up the underlying cluster hardware and the operating system are also discussed within this document. However, due to the fact that hardware platforms are evolving quickly, the paper can just give some guidelines on what hardware is required and how it should be configured. In addition, as Oracle Global Naming Service (GNS) is recommended for bigger clusters where nodes are dynamically joining and leaving, this kind of cluster configuration is not discussed in detail. Instead the paper focuses on cluster configurations running without GNS. On Microsoft Windows it may be difficult to integrate GNS into some customer´s existing network environments as it requires a DHCP server and DNS server on the same network.

Basic configuration steps on Oracle ASM will be described but more sophisticated Oracle ASM configurations like configurations with more than one storage system or storage based system copies are not covered within this document.

SAP specific configuration hints and recommendations will be provided in this document as well as suggestions about basic test scenarios.

Finally the document will provide information about SAP specific configuration steps in order to properly run BR*Tools and database tasks scheduled from SAP application servers.

Naming Conventions and Terms

Name or Synonym	Description / Purpose
Oracle ACFS	Oracle ASM Cluster File System is a cluster file system that uses Oracle ASM as the underlying volume manager infrastructure.
Oracle ASM	Oracle Automatic Storage Management is basically a volume manager for all Oracle Database related files and Oracle ACFS.
Oracle ASM Disk	An Oracle ASM disk is usually a physical disk with a logical partition on it. On Windows the logical partition is created as an extended partition with a logical volume in it.

Oracle ASM Disk Group	A number of Oracle ASM disks form an Oracle ASM disk group. An Oracle ASM disk group has 1 or more 'failure groups'. An extended partition can span part of a disk or a whole disk. Logical volumes are created inside an extended partition.
Failure Group	A named group to which an Oracle ASM disk belongs to. Oracle ASM guarantees to never store the same data twice within a failure group. Redundancy is achieved by distributing the same data to multiple failure groups.
OCR	Oracle Cluster Repository
RMAN	Oracle Recovery Manager
<DBNAME>	Specifies the name of the database (e.g. C11). <DBNAME> is derived from the Oracle initialization parameter 'db_name'. If 'db_unique_name' is set then <DBNAME> is derived from 'db_unique_name'.
<DBSID>	Specifies the name of the database instance. For single instance databases <DBSID> usually equals <DBNAME> e.g. C11. With RAC the number of the instance is appended e.g. C111, C112, etc. In some examples below the abbreviation <DBSIDn> will be used where 'n' equals the number of the RAC instance. The environment variable ORACLE_SID is usually set to <DBSID> or <ASMSID>. This the depends on the environment of the OS user
<SAPSID>	Specifies the name of the SAP system id. It is used for OS user names e.g. c11adm as well as for authentication of the OS user with Oracle e.g. OPS$C11ADM.
<IHRDBMS>	RDBMS Software Installation Home. The shared installation Oracle Home directory where the Oracle RDBMS Software is installed. This is a shared directory on Oracle ACFS.
<OHGRID>	The local path of the Oracle Home where the Grid Software was installed. E.g. 'C:\oracle\GRID\12102', 'C:\oracle\GRID\12102', … Note that for Grid Software Oracle Base it is required to use a local path outside of <OHGRID>. Therefore it is recommended to set Oracle Base to 'C:\oracle\BASE' during the Grid Software installation.

\<OHRDBMS\>	RDBMS Runtime Oracle Home. The runtime path of the Oracle Home where the RDBMS Software was installed. On Microsoft Windows this is the same as \<IHRDBMS\>. E.g. 'C:\oracle\\<DBNAME\>\12102', 'C:\oracle\\<DBNAME\>\12102'.
SCAN	The 'Single Client Access Name' is the name of a virtual host. In DNS the name should resolve to 3 different virtual IP addresses.
NODE-IP	The IP address of the node on the public network. The NODE-IP cannot failover and therefore identifies a physical cluster node.
SCAN-VIP	If SCAN was configured to resolve to 3 virtual IP addresses the SCAN-VIP is one of those. Please refer to http://www.oracle.com/technetwork/database/clustering/overview/scan-129069.pdf on Oracle TechNet for further information about SCAN.
NODE-VIP	The virtual IP address of a cluster node on the public network. The NODE-VIP can switch from one node to another if the owning node fails or is taken offline.
OIC	The 'Oracle Cluster Interconnect'. This is the private network of the cluster.

Installation and Migration Roadmaps

This chapter outlines the recommended installation path for new installations and upgrade path for upgrades from older Oracle single instance or clustered environments to Oracle RAC 12.1 with Oracle ASM. We recommend to always start with new hardware as this will leave your old system untouched and you will be able to abort the migration if unexpected problems arise. Steps that are common between the different installation paths are described in the chapter 'Common Setup and Installation Tasks'.

Note that migrations from Oracle 11.2 single instance to Oracle RAC 12.1 / Oracle ASM 12.1 will not be discussed here.

Each installation/upgrade path is based on the approach that an existing non-RAC Oracle database is being 'RAC enabled', moved to Oracle ASM and then put under the control of Oracle Clusterware. Therefore it is not recommended to upgrade Oracle Clusterware 11.2 to Oracle Clusterware 12.1. Instead the database will be upgraded to Oracle 12.1, enabled for RAC if necessary and finally added to the OCR of a fresh installed Oracle Clusterware 12.1.

Refer to the following SAP notes about Oracle 12.1 and upgrading to Oracle 12.1.

SAP Note #	Title
1914631	Oracle 12.1: Central Technical Note
1888485	Oracle 12.1.0.2: Database Parameter Settings
2229934	Oracle 12c RAC with Oracle ASM in Stretched Clusters
1915302	Oracle Database Software 12.1.0.2 Installation on MS Windows

New Installation

On Microsoft Windows Oracle ACFS is not supported for SAP installations. This means that it is not possible to install SAP Software on Oracle ACFS nor is it possible to install an Oracle RAC 12.1 system directly. If you would try to install SAP Software on Oracle ACFS SAPINST would fail when it detects that the Oracle RDBMS Home is on Oracle ACFS. Although new versions of SAPINST allow creating databases on Oracle ASM on Unix/Linux they do not offer this functionality on MS Windows and do not allow Oracle ACFS for a shared Oracle RDBMS Home.

Because of these restrictions we recommend to temporarily create local NTFS volumes, perform a complete standard database server installation ('Distributed Installation') with **local** Oracle RDBMS Home and initial database on NTFS and to move the database to Oracle ASM later when the installation is complete. In the case you want a shared Oracle RDBMS Home the local Oracle RDBMS Home(s) have to be deinstalled and reinstalled to ACFS mounted under the same directory path than the prior local Oracle RDBMS Home (e.g. "C:\ORACLE\<DBSID>" as the mountpoint and Oracle Base and "C:\ORACLE\<DBSID>\12102" as the Oracle RDBMS Home.)

Note that on MS Windows SAPINST creates new databases as single instance databases. Every fresh installed database is configured as a single instance database and requires 'RAC Enabling' which is discussed in a later chapter.

The following steps are required:

•   If you do not use new hardware remove any old operating systems and everything that is not needed. Get all the latest drivers for your hardware that might be required for Microsoft Windows 2012 R2.

•   Install and prepare the new operating system.

•   Install a local – temporary - Oracle RDBMS Home on NTFS.

•   Install an SAP database on NTFS if your SAPINST does not support Oracle ASM. Otherwise install the new SAP database on Oracle ASM.

•   Install and configure Oracle Grid Infrastructure Home.

•   Create all required Oracle ASM disk groups

- Install and configure Oracle RDBMS Home(s) on Oracle ACFS.

- Move your database to Oracle ASM with one of the Oracle Database migration / movement methods of your choice.

- 'RAC Enabling'.

- Install database server specific SAP Software (e.g. SAP BR*Tools).

Migrating from Oracle 12.1 Single Instance Database

Make sure you have a full backup of your database.

On new hardware

- Install and prepare the new operating system.

- Install and configure Oracle Grid Infrastructure Home and apply the latest SAP recommended bundle patch. Note that Oracle Grid Infrastructure 12.1 creates a mandatory 'management database' on Oracle ASM disk group OCR. For the required minimum size of the Oracle ASM disk group OCR refer to 'Oracle Grid Infrastructure Installation Guide for MS Windows'. We recommend 1.5 times of the minimal size.

- Create all required Oracle ASM disk groups

- Install and configure Oracle RDBMS Home(s) on Oracle ACFS. Ensure that your source system is on the same database release and patch level than you target system. E.g. 12.1.0.2 + latest recommended Windows bundle patch.

- Move your database to Oracle ASM with one of the Oracle Database migration / movement methods of your choice.

- 'RAC Enabling'.

- Install database server specific SAP Software (e.g. SAP BR*Tools).

On existing hardware

- If you do not use new 'plain' hardware remove any old operating systems and everything that is not needed but keep your NTFS volumes where your database and your old Oracle RDBMS Home resides. Get the latest drivers for your hardware that might be required for Microsoft Windows 2012 R2.

- Install and prepare the new operating system.

- Install and configure Oracle Grid Infrastructure Home and apply the latest SAP recommended bundle patch. Note that Oracle Grid Infrastructure 12.1 creates a mandatory 'management database' on Oracle ASM disk group OCR. For the required minimum size of the Oracle ASM disk group OCR refer to 'Oracle Grid Infrastructure Installation Guide for MS Windows'. We recommend 1.5 times of the minimal size.

- Create all required Oracle ASM disk groups

- Install and configure Oracle RDBMS Home(s) on Oracle ACFS. Ensure that your source system is on the same database release and patch level than your target system. E.g. 12.1.0.2 + latest recommended Windows bundle patch.

- Move your database to Oracle ASM with one of the Oracle Database migration / movement methods of your choice.

- 'RAC Enabling'.

- Install database server specific SAP Software (e.g. SAP BR*Tools).

Migrating from Oracle RAC 11.2

Assuming your database is already located on Oracle ASM it is strongly recommended to provide new hardware or at least a server where your existing Oracle ASM disk groups (or storage based copies of them) can be mounted to perform the database upgrade to Oracle 12.1.

Make sure you have a full backup of your database!

Migrating to new hardware

On your new hardware

- Install and prepare the new operating system.

- Install and configure Oracle Grid Infrastructure Home and apply the latest SAP recommended bundle patch. Note that Oracle Grid Infrastructure 12.1 creates a mandatory 'management database' on Oracle ASM disk group OCR. Therefore the space requirements for the Oracle ASM disk group OCR are significantly higher than with Oracle Grid Infrastructure 11.2. For the required minimum size of the Oracle ASM disk group OCR refer to 'Oracle Grid Infrastructure Installation Guide for MS Windows'. We recommend 1.5 times of the minimal size.

- Create all required Oracle ASM disk groups

On your old hardware:

- Shutdown each database instance using 'srvctl' and stop and disable your OracleService<DBSIDn> Service in Windows.

- Install an Oracle 12.1 RDBMS Home using 12.1.0.2 media and install the latest bundle patch released by SAP.

- Open a command prompt and set ORACLE_SID to your <DBSID> without the thread number at the end e.g. PRD not PRD1. Set ORACLE_HOME to the Oracle 12.1 home directory.

- Create a temporary OracleService<DBSID> Service in Windows using 'oradim.exe' from your Oracle 12.1 RDBMS Home.

- Prepare a database parameter file (pfile) for this database using the SAP recommended init.ora parameters for Oracle 12.1 RDBMS and your RAC specific parameters but with parameter 'cluster_database' set to FALSE.

- Startup your Oracle Database 11.2 in upgrade mode using the newly installed Oracle 12.1 RDBMS Home and with the newly created pfile to ensure the database opens successfully.
  E.g.
  ```
  sqlplus
  startup upgrade pfile=myupgradepfile.ora
  ```

- If the database instance starts, shutdown the database again and start a manual database upgrade as described in the Oracle documentation and the according SAP notes.

- Move your database to Oracle ASM with one of the Oracle Database migration / movement methods of your choice.

- 'RAC Enabling' - As the database was already RAC enabled not all steps are required.

- Install database server specific SAP Software (e.g. SAP BR*Tools).

On existing hardware

- Shutdown each database instance using 'srvctl' and stop and disable your OracleService<SIDn> Service in Windows.

- Install an Oracle 12.1 RDBMS Home using 12.1.0.2 media including the latest bundle patch released by SAP.

- Open a command prompt and set ORACLE_SID to your <SID> without the thread number at the end e.g. PRD not PRD1. Set ORACLE_HOME to the 12.1 home directory.

- Create a temporary OracleService<SID> Service in Windows using 'oradim.exe' from your newly installed Oracle 12.1 RDBMS Home.

- Prepare a pfile (e.g. myupgradepfile.ora) for this database using the SAP recommended init.ora parameters for Oracle 12.1 and your RAC specific parameters but with parameter 'cluster_database' set to FALSE. Keep the 'compatible' parameter at 11.2.0 until the whole upgrade process is completed and the ASM parameter 'compatible.rdbms' is raised to 12.1.0.

- Startup your Oracle 11.2 database in upgrade mode using the newly installed Oracle 12.1 RDBMS software and with the newly created pfile to ensure the database opens successfully.
  E.g.
  ```
  sqlplus
  startup upgrade pfile=myupgradepfile.ora
  ```

- If the database instance starts, shutdown the database again and start a manual upgrade (always run oradim from the path where it is installed!).

- Remove the Oracle 12.1 OracleService<SID> service with oradim.

- Setup your environment for Oracle 11.2, create a Oracle 11.2 OracleService<SID>, start the database and run the pre-upgrade script as described in the upgrade guide.

- Shutdown the database and remove the Oracle 11.2 OracleService<SID> service with oradim.

- Create an Oracle 12.1 OracleService<SID> service with oradim and perform a manual upgrade. Don´t forget to run SAP specific upgrade tasks and scripts.

- Backup the database.

- Remove your old Microsoft Windows installations from the cluster nodes.

- Install and prepare the new operating system.

- Install and configure Oracle Grid Infrastructure Home and apply the latest SAP recommended bundle patch. Note that Oracle Grid Infrastructure 12.1 creates a mandatory 'management database' on Oracle ASM disk group OCR. Therefore the space requirements for the Oracle ASM disk group OCR are significantly higher than with Oracle Grid Infrastructure 11.2. For the required minimum size of the Oracle ASM disk group OCR refer to 'Oracle Grid Infrastructure Installation Guide for MS Windows'. We recommend 1.5 times of the minimal size.

- Create all required Oracle ASM disk groups.

- Restore your database to a file system if required. If you have a copy of your database on a file system or network drive you can continue with moving the database to Oracle ASM.

- Move your database to Oracle ASM with one of the Oracle Database migration / movement methods of your choice.

Planning Oracle ASM Disk Groups

Shared disks are often located on an intelligent storage system. Although Oracle ASM is able and designed to work with thumb and cheap JBOD systems we strongly recommend using storage systems that already have built-in fault tolerance.

If your storage already provides fault tolerance we recommend creating all Oracle ASM disk groups with 'external redundancy' with the exception of the Oracle ASM disk group 'OCR' for which we recommend normal redundancy.

Some more advanced storage systems support additional capabilities like offlining and online disks or creating image copies of disks and therefore allow copying complete disk groups within the storage system which provides amazingly fast SAP system copies and disk image based backups. Both are discussed in detail in the whitepaper ' SAP with Oracle Real Application Clusters 11g Release 2 and Oracle Automatic Storage Management 11g Release 2 - Advanced Configurations & Techniques'.

If your storage does not provide any fault tolerance we recommend normal or high redundancy for all Oracle ASM disk groups. For Oracle ASM disk group 'OCR' we always recommend high redundancy.

When calculating space requirements and required number of disks it is always necessary to take the redundancy level and number of ASM disks of an Oracle ASM disk group into account.

OCR and Voting

Oracle Grid Infrastructure requires an Oracle ASM disk group for OCR and Voting Files. We strongly recommend creating an Oracle ASM disk group exclusively for OCR and Voting Files. For storage systems with built-in fault tolerance we recommend using 3 shared disks while JBOD systems should use 5 shared disks. All disks should be

of the same size. The minimum space required depends on the number of cluster nodes and the redundancy level of the Oracle ASM disk group. Please refer to 'Oracle Grid Infrastructure Installation Guide' for detailed information.

The recommended name for this Oracle ASM disk group is '**+OCR**' but you are free to choose another name.

Oracle ACFS

While Oracle Grid Infrastructure is installed on local disks of each of the cluster nodes Oracle RDBMS Software is installed on Oracle ACFS. Oracle ASM disk groups can contain multiple Oracle ACFS file systems the same way a standard disk can have multiple partitions each with its own file system.

We recommend creating one Oracle ASM disk group for all Oracle ACFS file systems and one Oracle ACFS file system per Oracle RDBMS Home. Keep in mind that since Oracle 11.2 patch sets are being installed out-of-place which means that you will need to reserve enough space for a second or third Oracle RDBMS Home.

The recommended size **for each** Oracle ACFS file system is at least 30GB for Oracle RDBMS Home, database logs and trace files from each database instance, BR*Tools logs, patch set installation.

The recommended minimum size of the Oracle ASM disk group that holds all the Oracle ACFS file systems therefore is <number of systems> * <size of Oracle ACFS file system>.

The recommended name for this Oracle ASM disk group is '**+ACFS**'.

E.g. for 3 SAP systems we recommend 3 Oracle ACFS file systems (each 30GB or more) on an Oracle ASM disk group with at least 60GB of free space.

Filesystemname	Mountpoint	Size [GB]
C11	C:\ORACLE\C11	30
C12	C:\ORACLE\C12	30
C13	C:\ORACLE\C13	30

Database Files

For an SAP installation on Oracle with Oracle ASM we recommend to choose one out of three variants on how to name and configure your Oracle ASM disk groups for database specific files. Which variant you choose depends on your backup strategy which in turn usually depends on the size of your database and required availability.

Variant 1 – small to medium data volumes, restore time not critical

Customer has small or medium sized databases where backup and/or restore + recovery of **all** databases can be accomplished by RMAN in a timely fashion.

Example: If a complete Oracle ASM disk group with data files from one or more databases is broken and all data files from all databases need to be restored to a newly created Oracle ASM disk group using RMAN.

Oracle ASM disk group recommendation:

Oracle ASM Disk Group Name	Stores
+DATA	- All data files - All temp files - Control file (first copy) - Online redo logs (first copy)
+ARCH	- Control file (second copy) - Archived redo logs
+RECO	- Control file (third copy) - Online redo logs (second copy) - RMAN backups (optional) - Fast recovery area (optional)

Variant 2 – medium to large data volumes, restore time important

Customer has databases of medium or large size where backup and/or restore + recovery of all databases cannot be accomplished in a timely fashion.

Usually customers will use RMAN and Split-Mirror techniques in combination.

Major differences to Variant 1 are:

- Separate Oracle ASM disk group for each database
- '<DBNAME>_' is the prefix for the name of the DATA disk group.

- Optionally a number is appended to the name of the DATA disk group if the database spans over more than one DATA disk group.

- Online redo logs are not stored in the 'DATA' disk groups. Instead online redologs are stored in an additional disk group that is used for the first member of each online redo log group.

Oracle ASM disk group recommendation:

Oracle ASM Disk Group Name	Stores
+<DBNAME>_DATA[#] e.g. ERP_DATA, CRM_DATA, ERP_DATA1 (e.g. for temporary additional space requirements during SAP upgrades or reorganizations)	- All data files - All temp files - Control file (first copy)
+OLOG	- Online redo logs (first copy)
+ARCH	- Control file (second copy) - Archived redo logs
+RECO	- Control file (third copy) - Online redo logs (second copy) - RMAN backups - Fast recovery area

Variant 3 – very large data and data change volumes, restore time crucial

Customer has very large databases where backup and/or restore + recovery of a single databases cannot be accomplished in a timely fashion.

Usually customers will use RMAN and Split-Mirror techniques in combination.

In this variant each relevant database file type is separated to different Oracle ASM disk groups.

Oracle ASM disk group recommendation:

Oracle ASM Disk Group Name	Stores

+<DBNAME>_DATA[#]
 	
e.g.   ERP_DATA,
 	
        CRM_DATA,
 	
        ERP_DATA1 (e.g. for temporary additional space requirements during SAP upgrades or reorganizations)
	• All data files
	• All temp files
	• Control file (first copy)

+<DBNAME>_ARCH e.g. ERP_ARCH, CRM_ARCH	• Control file (second copy) • Archived redo logs
+<DBNAME>_RECO e.g. ERP_RECO, CRM_RECO	• Control file (third copy) • RMAN backups • Fast recovery area
+<DBNAME>_OLOG e.g. ERP_OLOG, CRM_OLOG	• Online redo logs (first copy)
+<DBNAME>_MLOG e.g. ERP_MLOG, CRM_MLOG	• Online redo logs (second copy)

Important: Always place all files from one database into the according Oracle ASM disk groups of that database. Never place a file from one database in disk groups of another database.

Prerequisites for running Oracle Grid Infrastructure 12c Release 1 and Oracle Database 12c Release 1
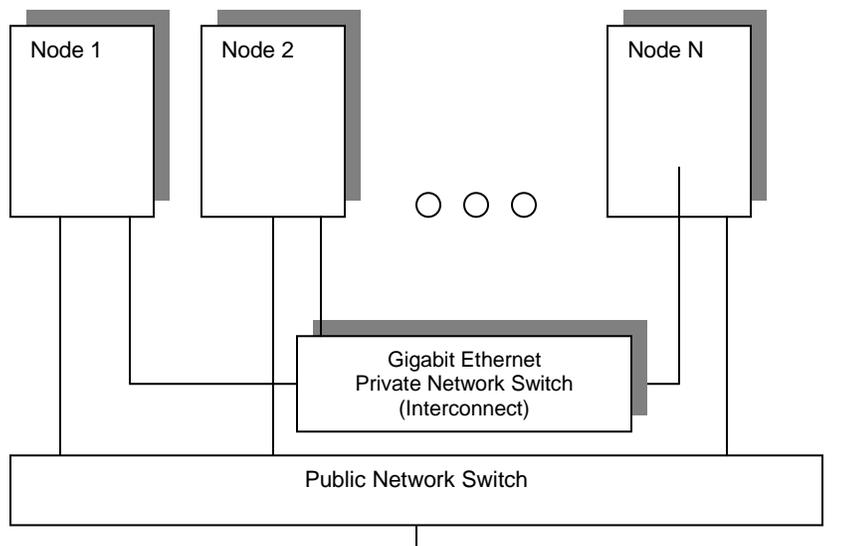
When starting with Oracle 12c RAC we recommend to study the hard- and software requirements discussed in the according Oracle documentation. We recommend using new cluster hardware as this allows leaving your old system untouched until the configuration of your new system is complete and your system is working properly.
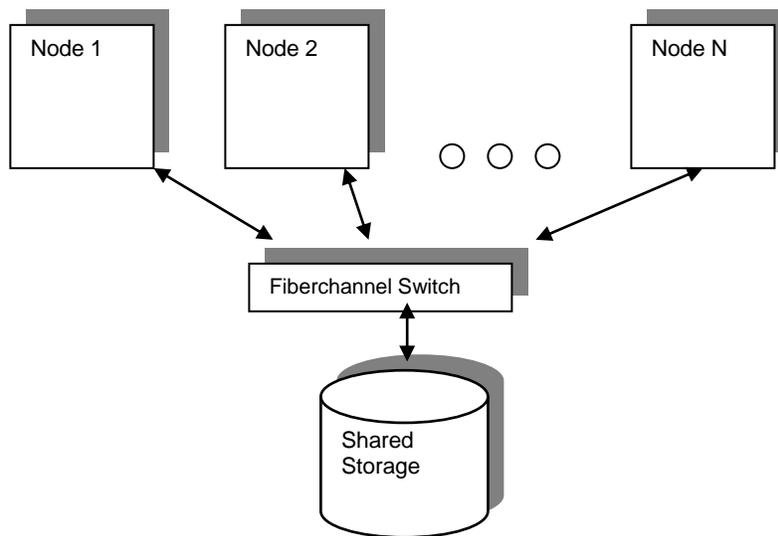
Hardware

Make sure your cluster is build from certified and well tested components. We recommend that each of the cluster nodes has the same number of CPU´s and network interfaces as well as the same amount of memory and local disk space. As Oracle Grid Infrastructure will be installed on a local volume make sure you have enough local disk space for installation and patching. The same is true for local Oracle RDBMS Home(s).

Note that patch sets for Oracle Grid Infrastructure and Oracle Database are installed out-of-place while patch set updates (or mini patches) are in-place.

Oracle Grid Infrastructure requires two network interfaces – one for the cluster interconnect and one for the public network. We recommend configuring 4 network interfaces into 2 pairs (teaming).

```
┌─────────┐  ┌─────────┐              ┌─────────┐
│ Node 1  │  │ Node 2  │    ○ ○ ○     │ Node N  │
│         │  │         │              │         │
└─────────┘  └─────────┘              └─────────┘

              ┌──────────────────────────┐
              │   Gigabit Ethernet        │
              │  Private Network Switch   │
              │     (Interconnect)        │
              └──────────────────────────┘
      ┌────────────────────────────────────────┐
      │        Public Network Switch            │
      └────────────────────────────────────────┘
```

Shared storage should be attached via fiber channel adaptors. For performance and availability reasons multiple fiber channel links and fiber channel switches between cluster nodes and storage system are recommended. This can be achieved by products like EMC Powerpath for example.

```
┌──────────┐   ┌──────────┐              ┌──────────┐
│  Node 1  │   │  Node 2  │   ○  ○  ○    │  Node N  │
└──────────┘   └──────────┘              └──────────┘
         ↘          ↕                  ↙
            ┌─────────────────────┐
            │ Fiberchannel Switch │
            └─────────────────────┘
                      ↕
                 ╭──────────╮
                 │  Shared  │
                 │ Storage  │
                 ╰──────────╯
```

Software

On each cluster node you need to be install the same operating system version and patches. We recommend using Microsoft Windows 2012 R2 Server 64Bit as this is the latest and most up-to-date Microsoft operating system.

Common Setup and Installation Tasks

Preparing Microsoft Windows 2012 for Oracle Grid Infrastructure

Before starting with Oracle Grid Infrastructure installation some mandatory pre-installation steps are required in order to be able to install Oracle Grid Infrastructure successfully.

Windows Registry Changes

Windows Time Synchronization

Ensure that Windows never automatically corrects system time backwards by setting MaxNegPhaseCorrection = 0 under 'HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\W32Time\Config' in the Registry.

See Microsoft Knowledge Base #884776.

Windows User Account Control and Automatic Elevation

•   Turn OFF Windows UAC:

```
Control Panel->System and Security->Change User Account Control Settings
```

•   Change automatic elevation settings:

```
UAC: Admin approval mode = Disabled
UAC: Behavior of the elevation prompt for administrators = Elevate without prompting
UAC: Run all administrators in Admin approval mode = Disabled
```
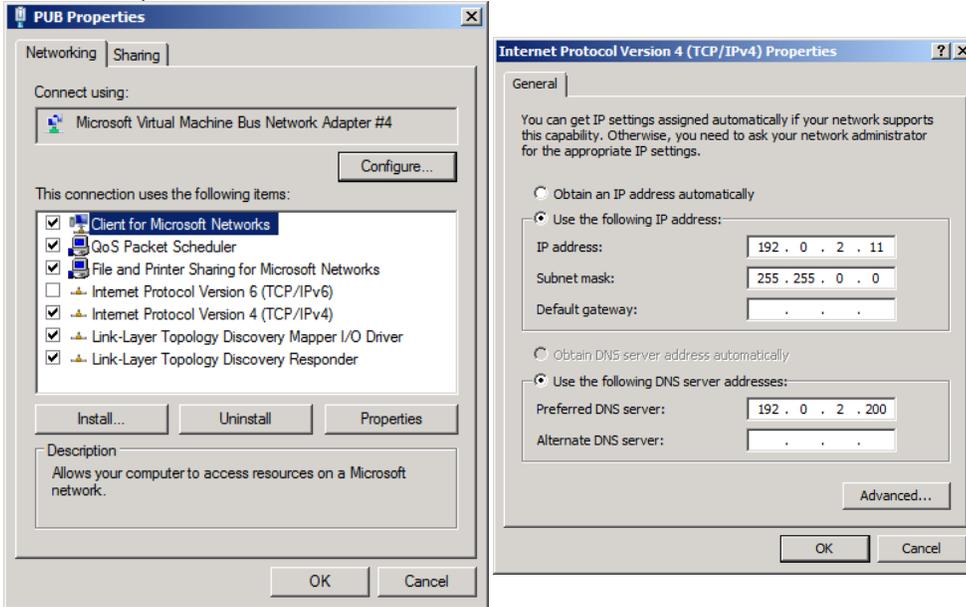
Network Infrastructure

Network Interfaces

•   Rename the network interfaces and choose speaking names for them. E.g. PUB, OIC.
Don´t use the name 'PUBLIC' as this will cause an error during the installation.

- Setup IP addresses, subnet mask and DNS of the network interfaces:

- Turn off IP V6 on all interfaces.

- Set the same bind order of all network interfaces. The public interface 'PUB' should be first, Oracle Interconnect 'OIC' second and then the other network interfaces (if any).
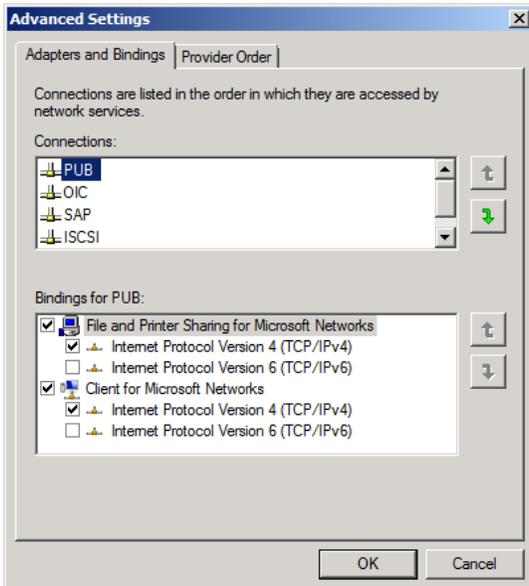
- Disable Windows Firewall for all network interfaces on all cluster nodes.

- Disable DHCP Media-Sensing (Microsoft Knowledge Base #239924)

- Turn OFF dynamic resizing of TCP/IP receive window (Microsoft Knowledge Base #934430) by issuing the following command from the command line prompt:

```
netsh interface tcp set global autotuning=disabled
```

- Query state by issuing:

```
netsh interface tcp show global
```

- Reboot and re-check everything. Also check if each NODE-IP and OIC-IP can be pinged.

DNS

Each cluster node requires 2 IP addresses and according hostnames on the public network: The node IP address and a node virtual IP address (NODE-VIP). In addition we recommend reserving 3 SCAN virtual IP addresses (SCAN-VIP). Create static entries in your DNS server for forward and reverse lookup.

E.g.

Forward-Lookup:

	oracx3	Host (A)	192.0.2.200	static
	oracx3	Host (A)	1.1.1.50	static
Forward Lookup Zones	oracx3	Host (A)	10.20.88.90	static
oracledev.wdf.sap.corp	oracx3v1	Host (A)	192.0.2.11	static
_msdcs	oracx3v1-vip	Host (A)	192.0.2.21	static
_sites	oracx3v2	Host (A)	192.0.2.12	static
_tcp	oracx3v2-vip	Host (A)	192.0.2.22	static
_udp	oracx3v3	Host (A)	192.0.2.13	static
DomainDnsZones	oracx3v3-vip	Host (A)	192.0.2.23	static
ForestDnsZones	scan	Host (A)	192.0.2.19	static
	scan	Host (A)	192.0.2.17	
	scan	Host (A)	192.0.2.18	

Reverse-Lookup:

	192.0.2.11	Pointer (PTR)	oracx3v1.oracledev.wdf.sa...	static
	192.0.2.12	Pointer (PTR)	oracx3v2.oracledev.wdf.sa...	static
	192.0.2.13	Pointer (PTR)	oracx3v3.oracledev.wdf.sa...	static
Reverse Lookup Zones	192.0.2.19	Pointer (PTR)	scan.oracledev.wdf.sap.corp.	static
2.0.192.in-addr.arpa	192.0.2.21	Pointer (PTR)	oracx3v1_vip.oracledev.wd...	static
	192.0.2.22	Pointer (PTR)	oracx3v2_vip.oracledev.wd...	static
	192.0.2.23	Pointer (PTR)	oracx3v3_vip.oracledev.wd...	static

For Oracle Cluster Interconnect (OIC) no entries are required in DNS but we recommend adding them to the 'hosts' file together with an OIC specific hostname.

E.g.

192.168.1.11   oracx3v1-oic
192.168.1.12   oracx3v2-oic
192.168.1.13   oracx3v3-oic

Windows Domain and OS Users

With Oracle RAC 11.2 a new user has been introduced that plays several roles:

•   Owner of all Oracle software components (Oracle Grid Infrastructure home and Oracle RDBMS home).

•   Administrator of Oracle Grid Infrastructure, Oracle ASM and databases.

The recommended name of this user is '**oracle**'.

Although the user '<sapsid>adm' still exists and is still required it is only used for running SAP software installed on the cluster nodes. E.g. for executing BR*Tools interactively or for running database related jobs from transaction DB13 via SAP Gateway Service.

•   Each cluster node should be member of your Windows Domain

•   Create user 'oracle' and user '<sapsid>adm' in your Windows domain.

•   Add user 'oracle' and '<sapsid>adm' to each cluster node's local 'Administrators' group. <u>It is not sufficient to add the user to a group in the Windows domain and to add the group to the local 'Administrators' group.</u>

With Oracle RAC 12.1 it is possible to run most Oracle-specific Windows Services either under the 'LocalSystem' account or with an unprivileged 'Oracle Home User' account. This paper focuses on the option using the 'LocalSystem' account. Please refer to SAP Note #1915302 for further information.

Oracle Administrator Groups

Create the following local groups on each cluster node and add your 'oracle' user as well as your '<sid>adm' user(s).

You can either create group ORA_DBA and ORA_OPER (recommended) or you can create groups ORA_<SID><instanceno>_DBA and ORA_<SID><instanceno>_OPER for each instance.

Preparing Oracle ASM Disks

When new disks are attached for the first time they are displayed as 'offline' and 'uninitialized'. Set them 'online' in the Windows disk administrator and initialize them. Before new disks can be used with Oracle ASM the disks must be prepared by creating an extended partition and a logical partition on it. This should be done with the 'diskpart' tool.

1)   select disk <#>

2)   clean

3)   create partition extended

4)   create partition logical

Important: When a new logical partition is created Microsoft Windows will ask to format the disk for use with a file system. Do not create a file system on those disks. Make sure that new state of the disk is visible to all nodes by initiating a scan of the SCSI bus or by rebooting the cluster nodes and remove all driveletters from those disks you plan to be Oracle ASM disks. This may require multiple scans or reboots if using a large number of disks. Make sure that nobody formats an Oracle ASM disk with a file system by documenting which disk is used for which purpose.

If you already have connected all shared disks you plan to use for your cluster prepare all the other disks in the same way.

E.g. After preparing 3 shared disks for Oracle Grid Infrastructure installation they should look like shown in the picture below.
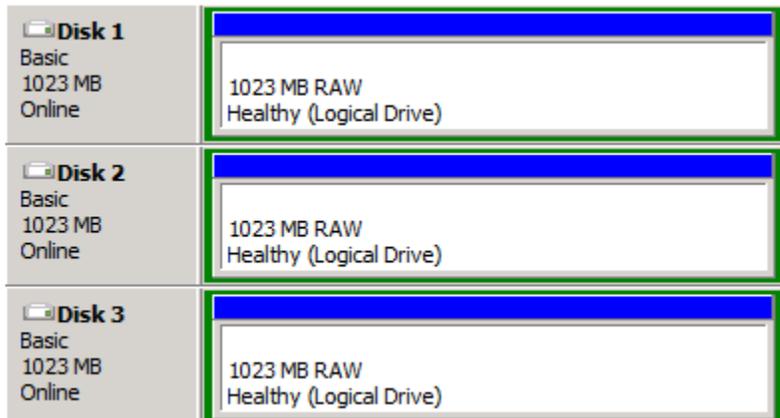
Oracle Grid Infrastructure Installation

As described above all installation and configuration steps regarding Oracle Software are performed as 'oracle' user. Make sure you are logged on as this user.

Setting environment variables as needed

The 'oracle' user is used on different cluster nodes with different Oracle Homes (OHGRID and one or more OHRDBMS'). Because of this we recommend to create small shell scripts to set the environment variables ORACLE_HOME and ORACLE_SID to the values needed. If you need to switch between the running instances you just need to run the according script.

E.g.

On cluster node oracx3v1:

```
C:\Users\oracle>type env_asm.cmd

set ORACLE_HOME=C:\ORACLE\GRID\12102

set ORACLE_SID=+ASM1


C:\Users\oracle>type env_c11.cmd

set ORACLE_HOME=C:\ORACLE\C11\12102

set ORACLE_SID=C111
```

On cluster node oracx3v2:

```
C:\Users\oracle>type env_asm.cmd

set ORACLE_HOME=C:\ORACLE\GRID\12102

set ORACLE_SID=+ASM2


C:\Users\oracle>type env_c11.cmd

set ORACLE_HOME=C:\ORACLE\C11\12102

set ORACLE_SID=C112
```

On cluster node oracx3v3:

```
C:\Users\oracle>type env_asm.cmd
```

```
set ORACLE_HOME=C:\ORACLE\GRID\12102

set ORACLE_SID=+ASM3


C:\Users\oracle>type env_c11.cmd

set ORACLE_HOME=C:\ORACLE\C11\12102

set ORACLE_SID=C113
```

Note that the instance number in ORACLE_SID does not necessarily match the number / order of the clusternode and may have to be adjusted.

E.g. OracleServiceASM2 could run on clusternode 3 (ORACLE_SID=+ASM2) while OracleServiceASM3 may run on clusternode 2 (ORACLE_SID=+ASM3).

These scripts will be used to set the environment of the 'oracle' user within the whole document.

Initial installation

Extract the installation media, open a command line prompt, move to that location and start 'setup.exe' with the following command line:

Screen 1: Choose the option to install and configure Oracle Grid Infrastructure for a cluster.

Screen 2: Choose the option for an advanced installation.

Screen 3: Select 'English' as the only product language.

Screen 4: Fill in the name of your cluster, the fully qualified DNS name of your SCAN-Listener, the port where the SCAN-Listener should listen (we recommend to keep port 1521) and do not configure a GNS.

Screen 5: Enter each node´s fully qualified hostname and virtual hostname.

Screen 6: Assign the role of each the network interface e.g.

Interfacename	Subnet	Usage
PUB	10.11.12.0	Public
OIC	192.168.1.0	Private
OTHERNETWORK	x.x.x.x	Do not use

Screen 7: Choose Standard-ASM (Flex-ASM is not supported with SAP).

Screen 8: Enter 'OCR' as the Oracle ASM disk group name, select the redundancy level of the disk group and click 'Stamp Disk'.

Screen 8a: In the popup dialog select 'Add or change label'.

Screen 8b: In the popup dialog select the disks you want to assign a label for and enter 'OCR' as the prefix for them. You can select multiple lines by holding down the <Shift> or <Ctrl> key.

Screen 8c: This popup dialog shows what label will be assigned to which disk. Click <Next>.

Screen 8d: Click <Finish>.

Screen 8: Your screen should now display the recently stamped disks as 'candidate disks'. Select them and click <Next>.

Screen 9: Choose 'Use same passwords for this account' enter the password and click <Next>.

Screen 10: Choose 'Do not use Intelligent Platform Management Interface (IPMI)' and click <Next>.

Screen 11: Check the checkbox if you want to register with Enterprise Manager Cloud Control otherwise leave it unchecked.

Screen 12: Select 'Predefined Windows-Account'.

Screen 13: Enter the directory for Oracle Base and Oracle Home. E.g. 'C:\ORACLE\BASE' for the Oracle Base of Oracle Grid Infrastructure and 'C:\ORACLE\GRID\12102' for the Oracle Home of Oracle Grid Infrastructure.

Start the installation by continuing through the next two screens (prerequisite check and installation summary).

After the installation and configuration of the Oracle Grid Infrastructure is finished check that Oracle Clusterware is running on all the cluster nodes.

```
crsctl status resource -t
```

If necessary upgrade the compatibility attributes 'COMPATIBLE.ASM', 'COMPATIBLE.RDBMS' and 'COMPATIBLE.ADVM' of disk group 'OCR' to '12.1.0.0.0'. Also choose a DISK_REPAIR_TIME that matches to your type of storage system. For storage systems with built-in fault tolerance you should select very high values (e.g. 240 hours). For storage systems without build-in fault tolerance you may want to keep the default value or even choose a smaller value.

E.g.

```
C:\Users\oracle>env_asm.cmd

C:\Users\oracle>set ORACLE_HOME=C:\ORACLE\GRID\12102

C:\Users\oracle>set ORACLE_SID=+ASM1

C:\Users\oracle>sqlplus

Enter user-name: / as sysasm

SQL> alter diskgroup OCR set attribute 'compatible.asm'='12.1.0.0.0';

Diskgroup altered.

SQL> alter diskgroup OCR set attribute 'compatible.rdbms'='12.1.0.0.0';

Diskgroup altered.

SQL> alter diskgroup OCR set attribute 'compatible.advm'='12.1.0.0.0';

Diskgroup altered.

SQL> alter diskgroup OCR set attribute 'disk_repair_time'='240h';

Diskgroup altered.
```

Creating Oracle ACFS File Systems

An Oracle ACFS file system is created in an Oracle ASM volume. An Oracle ASM volume is a specific Oracle ASM file type (like datafiles from a database) and is distributed over all Oracle ASM disks belonging to the same failure group. If the underlying Oracle ASM disk group is configured for normal or high redundancy Oracle ASM Dynamic Volume Manager (ADVM) - which sits between Oracle ACFS and ASM - ensures that the required number of redundant copies of a piece of data are written to Oracle ASM disks in the other failure groups.

To manage Oracle ASM disk groups and Oracle ACFS file systems we use the tool 'asmca'.

Creating a new 'Oracle ACFS' Disk Group

Open a new command prompt and set the environment to work with Oracle Grid Infrastructure and Oracle ASM: env_asm.cmd (created in an earlier chapter)

Screen 1: Shows a list of existing Oracle ASM disk groups. Click <Create>.

Screen 2: Click <Stamp Disks>.

Screen 3: Choose 'Add or change label' and click <Next>.

Screen 4: Choose the disks you want to stamp with an Oracle ASM link and choose a descriptive prefix for the labels e.g. 'Oracle ACFS'. Click <Next>.

Screen 5: This screen allows you the verify the name of the Oracle ASM link that will be written to each disk you selected. Click <Next> if everything is correct.

Screen 6: Click <Finish>

Screen 2: Now you are back on screen 2. Enter a name for the new Oracle ASM disk group (e.g. 'ACFS'), select the redundancy level and select the disks you want to add to the Oracle ASM disk group. Click on 'Show Advanced Options' and enter '12.1.0.0.0' (without quotes) for ASM-, Database- and ADVM compatibility. Finally click <OK> to create the new Oracle ASM disk group.

Use this procedure to create all your Oracle ASM disk groups.

Always set the compatibility attributes to '12.1.0.0.0' or higher as this will enable some functionalities you may want to use later.

| Oracle ASM Disk Group Name	Size [GB]	Free [GB]	Usable [GB]	Redundancy	State
ARCH	9.76	9.63	9.63	EXTERN	MOUNTED(3of3)
ACFS	24.41	24.28	24.28	EXTERN	MOUNTED(3of3)
DATA	39.06	39.06	38.93	EXTERN	MOUNTED(3of3)
RECO	39.06	39.06	38.93	EXTERN	MOUNTED(3of3)
OCR	30.02	15.01	7.45	NORMAL	MOUNTED(3of3)

Creating the Oracle ACFS Volume

Although it is possible to create just one Oracle ACFS volume and file system for all your Oracle RDBMS Homes we recommend creating one Oracle ACFS volume and file system for each Oracle RDBMS Home for maintenance reasons.

Screen 1: Click the 'Volumes' tab.

Screen 2: This screen shows all existing Oracle ACFS volumes. Click <Create>.

Screen 3: Choose a volume name, select the appropriate Oracle ASM disk group, enter the required size of the volume and click <OK>.
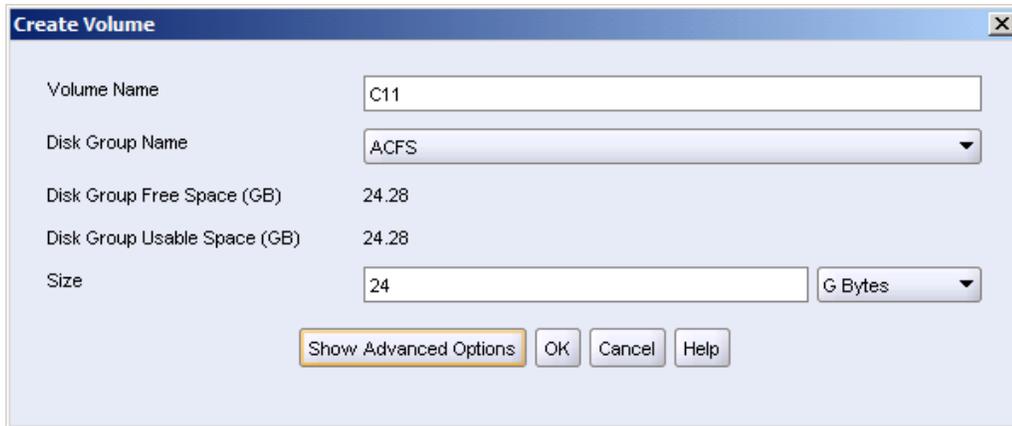
Screen 2: Back on screen 2 you will see the newly created Oracle ASM volume.

It is very important to make a note of the name of the volume (e.g. C11) together with the device name (e.g. '\\.\asm-c11-210'). You need the device name when doing ACFS filesystem checks at a later time. Oracle ACFS checks as well as Oracle ACFS backups and restores are discussed in chapter Working with Oracle ACFS.

Creating an Oracle ACFS

Screen 1: This screen show a list of all existing Oracle ASM. Click <Create>

Screen 2: Choose the ACFS volume where you want to create the ACFS on, select 'Database Home File System', enter the mountpoint and click <OK>.

E.g. Choose 'C11 \\.\asm-c11-210' as the ACFS volume and 'C:\ORACLE\C11' as the database home mount point.

The newly created ACFS should now be listed in the filesystems tab of 'asmca'.

Check that all cluster nodes have mounted the new file system.

RDBMS Home Installation

The installation of the Oracle RDBMS Home(s) is very similar to installations in a non-clustered environment. The major difference is that during the installation in a cluster each node´s registry is updated and a central inventory is created on each node.

Extract the installation media to a staging directory and open a command prompt. Move to the staging directory and run setup.exe

Screen 1: Uncheck 'I wish to receive security updates via My Oracle Support'. You may receive the warning that you did not specify your email address.

Screen 2: Choose 'Install database software only'.

Screen 3: Choose 'Oracle Real Application Clusters database installation'.

Screen 4: Make sure each cluster node is checked.

Screen 5: Make sure 'English' is the only selected language.

Screen 6: Choose 'Enterprise Edition'.

Screen 7: Choose 'Predefined Windows-Account'.

Screen 8: Enter 'Oracle Base' and 'Software Location'.

Note that in this dialog Oracle Base is different from Oracle Base used during the installation of the Oracle Grid Infrastructure. Both directories have to be on Oracle ACFS.

E.g. 'C:\ORACLE\C11' for Oracle Base and 'C:\ORACLE\C11\12102' for Oracle Home (Software-Location).

Start installation by continuing to the prerequisite check and installation overview.

Repeat the installation procedure for each additional Oracle RDBMS Home you want to install.

Patching Oracle Grid Home and Oracle RDBMS Homes

How to install patches is well documented in the according patch notes and readme. Although, based on our experience, we recommend some additional steps to avoid problems like open files that cannot be overwritten or cluster resources in transition during patching.

We strictly recommend to follow the Non RAC rolling installation.

Execute each step on each cluster node and wait for completion at each step:

E.g.

Execute the required commands of step 1 on node A and wait for completion of each command.

Execute the required commands of step 1 on node B and wait for completion of each command.

Execute the required commands of step 1 on node C and wait for completion of each command.

.

.

Execute the required commands of step 7 on node A and wait for completion of each command.

Execute the required commands of step 7 on node B and wait for completion of each command.

Execute the required commands of step 7 on node C and wait for completion of each command.

At this point every Oracle specific file that may be overwritten during patch operation should be closed.

In the case opatch finds an open file use the command 'tasklist /m' or ProcessExplorer to find the according process.

Check if you need to install a new version of opatch into each Oracle Home before you start patching.

To patch your Grid Home open a new command prompt and set ORACLE_HOME to the Grid Home (e.g. C:\ORACLE\GRID\12102). If you have created shell scripts for switching environments just run the script to set your 'GRID/ASM' environment.

Follow the instructions in the patch readme in order to patch your Oracle RDBMS Home(s).

RAC Enabling

Database specific changes

All database specific changes can be performed from one cluster node.

Note that with Oracle ASM filenames are generated by 'Oracle Managed Files for ASM'. The parameters 'DB_CREATE_FILE_DEST', 'DB_CREATE_ONLINE_LOG_DEST_n', 'DB_RECOVERY_FILE_DEST' specify the Oracle ASM disk group where Oracle will create the files by default. This means that there is no need to specify the filenames for datafiles – OMF will generate a name based on the tablespace name. The names for the controlfiles of the database must be specified as Oracle ASM aliases in the spfile. In the case you need to restore controlfiles there is no need to change the names in the spfile.

Sometimes it makes sense to specify a full filename on Oracle ASM. E.g. if you have manually copied a database from NTFS to Oracle ASM using the 'asmcmd' utility you may want to use aliases for the destination filename in order to create a new controlfile after each database file has been moved to Oracle ASM.

Undo Tablespaces

Ensure that each Oracle RDBMS instance has its own dedicated undo tablespace of the appropriate size.

If you have upgraded from Oracle RAC 11.2 and the number of cluster nodes did not change, you should already have undo tablespaces for all instances and you can skip this step because the undo tablespaces will be available after moving the database to Oracle ASM.

E.g.
```
create undo tablespace PSAPUNDO1 datafile size 32000M autoextend on;
create undo tablespace PSAPUNDO2 datafile size 32000M autoextend on;
```

Online Redologs

We recommend using 'DB_CREATE_ONLINE_LOG_DEST_n' to set the default Oracle ASM disk group for online redologs and their duplexed mirror.

E.g. If you decided to create dedicated Oracle ASM disk groups for online redologs '+C11OLOG' (OLOG for origlog) and '+C11MLOG' (MLOG for origlog) you should set to

```
*.DB_CREATE_ONLINE_LOG_DEST_1='+C11OLOG'
*.DB_CREATE_ONLINE_LOG_DEST_2='+C11MLOG'
```

and create new online redolog groups after moving your database to Oracle ASM.

Each RAC instance uses a specific thread for writing redo logs. You therefore need to create new redologs.

E.g.
```
alter database add logfile thread 1 group 10 size 250M;

alter database add logfile thread 1 group 11 size 250M;
alter database add logfile thread 1 group 12 size 250M;
alter database add logfile thread 2 group 20 size 250M;
alter database add logfile thread 2 group 21 size 250M;
alter database add logfile thread 2 group 22 size 250M;
.
.
```

You should add the same number of redo log groups as in a single node environment. The redo log sizes should also be the same.

If you have plenty of disk space on your Oracle ASM disk group, it's a good idea to use this additional disk space for more online redologs. This will reduce the chance that the system requires archived redologs during crash recovery.

For all new redolog groups that have been added, the corresponding thread has to be enabled:

```
alter database enable thread 2;
```

SPFILE

It is recommended and SAP standard to place the spfile that belongs to a database on Oracle ASM. SAP BR*Tools will automatically create backups from the spfile on Oracle ASM and allow to modify initialization parameters directly.

Starting with Oracle Grid Infrastructure 12c we strictly do not recommend to place the spfile in 'OHRDBMS\dbs' as the semantics of the create spfile command did change and will always update the according property of the database resource in OCR with the given location. It will not check if this location is accessible from each clusternode.

Assuming that you have a pfile with all the initialization parameters create a spfile on Oracle ASM from it.

E.g.
```
create spfile='+DATA/C11/spfile.ora' from pfile='C:\temp\finalpfile.ora';
```

This command creates an Oracle ASM alias that points to the active spfile on Oracle ASM. The alias is required because Oracle ASM will use an OMF name for the newly create spfile and we want to make sure that the spfile is always found under the same name. This is important when we add the database and instances to the cluster repository.

RAC specific parameters are:

```
*.cluster_database=TRUE

*.undo_management=auto

<sid><instancenumber>.instancenumber=<instancenumber>

<sid><instancenumber>.thread=<threadnumber>

<sid><instancenumber>.undo_tablespace='<nameofundotablespace>'

<sid><instancenumber>.remote_listener='<scan:port>'
```

Oracle ASM specific parameters are:

```
*.DB_CREATE_ONLINE_LOG_DEST_1='<+DISKGROUP>'

*.DB_CREATE_ONLINE_LOG_DEST_2='<+DISKGROUP>'

*.DB_CREATE_FILE_DEST='<+DISKGROUP>

*.log_archive_dest_1='location=<+DISKGROUP>/<DBNAME>/ORAARCH
valid_for=(ALL_LOGFILES,ALL_ROLES) db_unique_name=<DBNAME>'

*.control_files='<+DISKGROUP1>/<DBNAME>/CNTRL<DBNAME>.DBF',

'<+DISKGROUP2>/<DBNAME>/CNTRL<DBNAME>.DBF',

'<+DISKGROUP3>/<DBNAME>/CNTRL<DBNAME>.DBF'
```

Optional:

```
*.DB_RECOVERY_FILE_DEST='<+DISKGROUP>'

*.DB_RECOVERY_FILE_DEST_SIZE=<size>G
```

The following example shows a full working example of a real world production system with 3 instances:

```
C111.__oracle_base='c:\oracle\C11'#ORACLE_BASE set from environment

C112.__oracle_base='c:\oracle\C11'#ORACLE_BASE set from environment

C113.__oracle_base='c:\oracle\C11'#ORACLE_BASE set from environment

*._awr_mmon_deep_purge_all_expired=TRUE#SAP RECOMMENDED SETTINGS

*._fix_control='20355502:8','5099019:ON','5705630:ON','6055658:OFF','6120483:OFF','63
99597:ON','6430500:ON','6440977:ON','6626018:ON','6972291:ON','7168184:OFF','7658097:
ON','8937971:ON','9196440:ON','9495669:ON','13077335:ON','13627489:ON','14255600:ON',
'14595273:ON','18405517:2'#SAP RECOMMENDED SETTINGS

*._ktb_debug_flags=8#SAP RECOMMENDED SETTINGS

*._mutex_wait_scheme=1#SAP RECOMMENDED SETTINGS

*._mutex_wait_time=10#SAP RECOMMENDED SETTINGS

*._optim_peek_user_binds=FALSE#SAP RECOMMENDED SETTINGS

*._optimizer_adaptive_cursor_sharing=FALSE#SAP RECOMMENDED SETTINGS

*._optimizer_aggr_groupby_elim=FALSE#SAP RECOMMENDED SETTINGS

*._optimizer_batch_table_access_by_rowid=FALSE
```

```
*._optimizer_extended_cursor_sharing_rel='NONE'#SAP RECOMMENDED SETTINGS

*._optimizer_reduce_groupby_key=FALSE

*._optimizer_use_feedback=FALSE#SAP RECOMMENDED SETTINGS

*._securefiles_concurrency_estimate=50#SAP RECOMMENDED SETTINGS

*._suppress_identifiers_on_dupkey=TRUE#SAP RECOMMENDED SETTINGS

*._use_single_log_writer='TRUE'#SAP RECOMMENDED SETTINGS

*.archive_lag_target=3600

*.AUDIT_FILE_DEST='C:\Oracle\C11\saptrace\audit'

*.compatible='12.1.0.2.0'

*.cluster_database=true

*.cluster_database_instances=3

*.control_file_record_keep_time=30

*.control_files='+C11_DATA/C11/CNTRLC11.DBF','+C11_ARCH/C11/CNTRLC11.DBF','+C11_RECO/
C11/CNTRLC11.DBF'

*.db_block_size=8192

*.db_cache_size=15G

*.db_create_file_dest='+C11_DATA'

*.db_create_online_log_dest_1='+C11_OLOG'

*.db_create_online_log_dest_2='+C11_MLOG'

*.db_flashback_retention_target=20160

*.db_files=2500

*.db_name='C11'

*.DB_RECOVERY_FILE_DEST='+C11_RECO'

*.DB_RECOVERY_FILE_DEST_SIZE=950G

*.DIAGNOSTIC_DEST='C:\Oracle\C11\saptrace'

*.event='10027','10028','10142','10183','10191','10995 level 2','38068 level
100','38085','38087','44951 level 1024'#SAP RECOMMENDED SETTINGS

*.FILESYSTEMIO_OPTIONS='setall'

C111.instance_number=1

C112.instance_number=2

C113.instance_number=3
```

```
*.log_archive_dest_1='location=+C11_ARCH/C11/ORAARCH
valid_for=(ALL_LOGFILES,ALL_ROLES) db_unique_name=C11'

*.log_archive_format='C11ARC%t_%s_%r.dbf'

*.log_checkpoints_to_alert=true

*.max_dump_file_size='20000'

*.open_cursors=2000

*.optimizer_adaptive_features=FALSE

*.optimizer_capture_sql_plan_baselines=FALSE

*.optimizer_index_cost_adj=20

*.parallel_execution_message_size=16384

*.parallel_max_servers=160#DB-CPU-Cores * 10    SQL> show parameter cpu_count

*.parallel_threads_per_cpu=1

*.pga_aggregate_target=10G#OLTP: 20 % of available memory

*.processes=1000#(#ABAP-Workprozesse * 2) + (#J2EE server processes * <max.
Verbindungen> ) + PARALLEL_MAX_SERVERS + 40

*.query_rewrite_enabled='false'

*.recyclebin='off'

*.replication_dependency_tracking=false

*.remote_listener='scanwp.sap.swm.de:1521'

*.sessions=2000# 2 * <PROZESSE>

*.shared_pool_size=4G# 400 MB or more (SAP-Note 690241)

*.undo_retention=900

C111.thread=1

C112.thread=2

C113.thread=3

C111.undo_tablespace='PSAPUNDO'

C112.undo_tablespace='PSAPUNDO2'

C113.undo_tablespace='PSAPUNDO3'
```

Please note that it is not possible to start the database until "Oracle Administrator Groups" and the Windows services for the instances have been created and the database and it's instances have been configured with Oracle Clusterware.

SQL*Net Configuration

Note that with Oracle Grid Infrastructure 12.1 no manual configuration of listener.ora and sqlnet.ora is required on server side except of the SID_LIST_LISTENER entry that is required for BR*Tools to startup/shutdown the database instances from remote. The installation process creates the files listener.ora, sqlnet.ora and endpoints_listener.ora. Oracle Clusterware will automatically updates the listener.ora and endpoint_listener.ora files based on the current status of the listener resources and the database services that have been defined.

With Oracle 12.1 RAC listener.ora, sqlnet.ora and endpoints_listener.ora are located under 'OHGRID\network\admin' and not in one of the Oracle RDBMS Homes. Note that this is a local path and that you have to modify the listener.ora on each cluster node.

E.g. On Node A:

```
SID_LIST_LISTENER=
(
  SID_LIST=
    (SID_DESC=(SID_NAME=C111)(ORACLE_HOME=C:\ORACLE\C11\12102))
)
```

E.g. On Node B:

```
SID_LIST_LISTENER=
(
  SID_LIST=
    (SID_DESC=(SID_NAME=C112)(ORACLE_HOME=C:\ORACLE\C11\12102))
)
```

On the client side the SAP application servers do not require an address list of all the listeners on each node. Instead each client needs to know to which database service to connect to as well as the name and port of the SCAN. This can be done in two ways: Either by creating a tnsnames.ora or by using the EZCONNECT method (recommended).

When using EZCONNECT set the environment variable dbs_ora_tnsnames to '//<hostname>:<port>/<servicename>.

If you still want to use tnsnames.ora please see the example below:

tnsnames.ora

```
C11.WORLD=
  (DESCRIPTION =
    (ADDRESS_LIST =
        (ADDRESS =
           (COMMUNITY = SAP.WORLD)
           (PROTOCOL = TCP)
           (HOST = scan.oracledev.wdf.sap.corp)
           (PORT = 1521)
        )
    )
    (CONNECT_DATA =
        (SERVICE_NAME = cic11)
    )
  )
```

Oracle Instance Creation

When creating or removing an Oracle instance service (OracleService<SID>) using oradim make sure you always run oradim from the correct Oracle RDBMS Home and from the cluster node where you want to create the new service.

You should never rely on the PATH variable when running oradim.

To create a new instance service:
*   open a new command prompt
*   set ORACLE_HOME=<OHRDBMS> e.g. C:\ORACLE\C11\12102
*   cd %ORACLE_HOME%\bin
*   oradim -new -sid <DBSIDn> -startmode=manual

To remove an instance service:
*   open a new command prompt
*   set ORACLE_HOME=<OHRDBMS> e.g. C:\ORACLE\C11\12102
*   cd %ORACLE_HOME%\bin
*   oradim -delete -sid <dbsid+instanceno>

E.g. for a database 'C11' that should have instances on 3 cluster nodes you would create 3 new instance services:

```
Cluster node 1: oradim –new –sid C111 –startmode=manual
Cluster node 2: oradim –new –sid C112 –startmode=manual
Cluster node 3: oradim –new –sid C113 –startmode=manual
```

Create the Oracle Clusterware Resources

Setup your environment to work with Oracle Grid Infrastructure by calling your shell script that sets up the environment accordingly. E.g. 'env_asm.cmd'.

Add the database. E.g.

```
srvctl add database -d C11 -o C:\ORACLE\C11\12102 -p +C11_DATA/C11/SPFILEC11.ORA
```

Add instances that belong to the database. E.g.

```
srvctl add instance -d C11 -i C111 -n oracx3v1
srvctl add instance -d C11 -i C112 -n oracx3v2
srvctl add instance -d C11 -i C113 -n oracx3v3
```

Add database services that should be managed by Oracle Clusterware.

Each Oracle RAC instance must have a service named like to instance itself and the service must be configured to run on this instance only.

It is required to create one database service for your SAP central instance and for each SAP application server you require additional database services (i.e. for n SAP application servers you need n database services where every SAP application server is connected to its own database service) – not counting the instance specific services.

Never configure a database service to run on more than one database instance at the same time. This would create connections from one application server to more than one database instance, which may cause serious performance and availability problems.

E.g.

```
srvctl add service -d C11 -s C111 -r "C111"

srvctl add service -d C11 -s C112 -r "C112"

srvctl add service -d C11 -s C113 -r "C113"


srvctl add service -d C11 -s C11CI -r "C111" -a "C112,C113"

srvctl add service -d C11 -s C11AS1 -r "C112" -a "C113,C111"

srvctl add service -d C11 -s C11AS2 -r "C113" -a "C111,C112"
```

Start the database services:

srvctl start service –d C11 (will try to start all services defined for database 'C11')

SAP-specific database changes

DDLOG Sequence

SAP is using a database sequence called "DDLOG" to generate timestamps for various purposes. Therefore SAP needs consecutive numbers returned from DDLOG. To guarantee this in a RAC environment the sequence has to be recreated in "ordered" mode. Please refer to SAP note #743555 for a description on how to do this.

Checking the Setup

In order to check the status of the cluster´s local resources and cluster resources setup the appropriate environment in your command line and run 'crsctl status resource –t' to print a complete list. Note that every resource except 'ora.gsd' and 'ora.oc4j' should have the status shown as 'TARGET ONLINE STATUS ONLINE'.

•   Shutdown and power off each cluster node and then reboot them to see if every node comes up properly and all the cluster services are correctly restarted. Note that it can take ~15 minutes until all cluster resource are started.

•   Use tnsping to verify that all listeners are reachable and the according services are registered properly.

•   Try to connect to each database service.

•   Check if database instances can be started and stopped as expected using 'srvctl' utility.

•   Check the alert logs for errors.

Reusing Oracle ASM Disks

If a disk is added to an Oracle ASM disk group for the first time the disk gets flagged as 'in use'. This means that even if you remove this Oracle ASM disk from the Oracle ASM disk group you may not be able to reuse it in the same or another Oracle ASM disk group. This is also the case if you decide to re-install Oracle Grid Infrastructure using disks for OCR that have been used earlier.

The easiest way to solve this problem is to format the disk with NTFS and then use the 'diskpart' utility to remove all partitions and recreate them. After this 'asmca' can be used to put a new label on the disk.

Do not forget to issue a 'rescan disks' on the other cluster nodes and to remove driveletters that may have been assigned by Microsoft Windows.

Working with Oracle ACFS

Checking Oracle ACFS

Oracle Grid Infrastructure 12.1 includes a utility to perform offline file system checks. To unmount an mounted Oracle ACFS volume all processes with open handles on that file system must be stopped.

This includes the services 'OracleService<SIDn>' and 'Oracle <DBSIDn> VSS Writer Service' and all other programs (command prompts, tail, etc.) that may have open handles on that file system on all cluster node.

- Shutdown your database that is run out of the Oracle 12.1 RDBMS Home you want to check.
E.g. `srvctl stop database –d <DBSID>`

- Retrieve the virtual device name of the Oracle ASM volume housing the Oracle ACFS file system.
E.g. acfsutil info fs or the names of the virtual Oracle ASM volumes you wrote down during the installation.

- Stop the file system and wait for unmount to complete.
E.g. `srvctl stop file system -d \\.\asm-c11-210`

If an error is thrown the reason can be one of the following:
1) Unmounting the file system was too slow. Run 'srvctl stop file system...' command again.
2) You still have some files open. Check whether a Windows service is still running or whether another program is still accessing the file system. E.g. A command prompt where you did a 'cd' to a directory on Oracle ACFS requires a handle. So either 'cd' out of the directory or close the command prompt.

E.g.
`C:\Users\oracle>srvctl stop database -d C11`

```
C:\Users\oracle>acfsutil info fs
C:\ORACLE\C11
    ACFS Version: 12.1.0.2.0
    flags:         MountPoint,Available
    mount time:    Wed Aug 17 15:33:38 2015
    volumes:       1
    total size:    25769803776
    total free:    21699956736
    primary volume: asm-c11-210
        label:
        flags:                 Primary,Available,ADVM
        on-disk version:       39.0
        allocation unit:       4096
        size:                  25769803776
        free:                  21699956736
        ADVM diskgroup         ACFS
        ADVM resize increment: 268435456
        ADVM redundancy:       unprotected
        ADVM stripe columns:   4
        ADVM stripe width:     131072
    number of snapshots:  0
    snapshot space usage: 0


C:\Users\oracle>srvctl stop file system -d \\.\asm-c11-210
PRCR-1065 : Failed to stop resource ora.acfs.c11.acfs
CRS-5014: Agent "C:\ORACLE\GRID\12102\bin\orarootagent.exe" timed out starting
process "C:\ORACLE\GRID\12102\bin\acfssinglefsmount.bat" for action "stop": details
at "(:CLSN00009:)" in
"C:\ORACLE\GRID\12102\log\oracx3v1\agent\crsd\orarootagent\orarootagent.log" CRS-
5017: The resource action "ora.acfs.c11.acfs stop" encountered the following error:
(:CLSN00009:)Utils:execCmd aborted

CRS-2675: Stop of 'ora.acfs.c11.acfs' on 'oracx3v1' failed
CRS-2678: 'ora.acfs.c11.acfs' on 'oracx3v1' has experienced an unrecoverable failure
```

Trying again...

```
C:\Users\oracle>srvctl stop file system -d \\.\asm-c11-210
PRCC-1016 : acfs.c11 was already stopped
PRCR-1005 : Resource ora.acfs.c11.acfs is already stopped


C:\Users\oracle>
```

• If the file system was unmounted successfully check it with the 'acfschkdsk' utility. Calling acfschkdsk without any parameters displays a list of options.

```
C:\Users\oracle>acfschkdsk \\.\asm-c11-210
version                  = 12.1.0.2.0
*****************************
********** Pass: 1 **********
*****************************
Oracle ASM Cluster File System (ACFS) On-Disk Structure Version: 39.0
ACFS file system created at: Thu Aug 11 13:33:05 2011
checking primary file system
       Files checked in primary file system: 85%
       Files checked in primary file system: 100%
        Checking if any files are orphaned...
        0 orphans found
Checker completed with no errors.
```

- Restart the file system.

```
C:\Users\oracle>srvctl start file system -d \\.\asm-c11-210
```

- Wait for 3-4 minutes and check if each cluster node did mount the file system in the Windows Explorer before you restart your database instances.

Backup and Restore

Oracle ACFS does not support backup tools that rely on Microsoft Windows Volume Shadow Copy Services. Other backup tools that do not require VSS should work fine with Oracle ACFS. We recommend to create a full backup of an Oracle ACFS volume (which basically just holds Oracle 12.1 RDBMS Homes, Oracle ADR specific files and logs generated by BR*Tools) after initial installation and configuration and after each modification of the installed software or of the configuration files (e.g. after patch installations).

Before you start patching Oracle software located on Oracle ACFS you must shutdown the database that is run from that Oracle RDBMS Home and to unmount and check the file system. If the file system has unrecoverable errors – which is very unlikely - we recommend to recreate it and to restore the latest backup.

Configuring BR*Tools for Oracle ASM

SAP has extended the functionality of BR*Tools to support Oracle ASM. Place the init<sid>.sap configuration files for archivelog backup and database backup in the 'sapprof' directory. Note that you need specific configuration files for each cluster node on which you run backups as the init<sid>.sap configuration files need different values for parameter 'asm_ora_sid' on each node.

Example init<sid>.sap files:

Archivelog backup:

```
backup_mode = all

restore_mode = all

backup_type = online

backup_dev_type = rman_disk

backup_root_dir = C:\oracle\C11\sapbackup

stage_root_dir = C:\oracle\C11\sapbackup

compress = no

compress_cmd = "compress -c $ > $"

uncompress_cmd = "uncompress -c $ > $"

compress_dir = C:\oracle\C11\sapreorg

archive_function = save_delete

archive_copy_dir = C:\oracle\C11\sapbackup

archive_stage_dir = C:\oracle\C11\sapbackup

tape_copy_cmd = cpio

disk_copy_cmd = copy

stage_copy_cmd = rcp

pipe_copy_cmd = rsh

cpio_flags = -ovB

cpio_in_flags = -iuvB

cpio_disk_flags = -pdcu

rman_channels = 1

rman_filesperset = 1

_rman_sess_stmt = ("alter session set optimizer_mode=RULE")

rman_send = (

"channel sbt_1 'NSR_ENV=(NSR_SERVER=xxbackapp02-
se,NSR_CLIENT=oracx3v1,NSR_DATA_VOLUME_POOL=IDES Oracle 4
Weeks,NSR_GROUP=ArchiveLogs_C11)'"

)

rewind = "mt -f $ rewind"

rewind_offline = "mt -f $ offline"

tape_pos_cmd = "mt -f $ fsf $"

tape_size = 100G
```

```
exec_parallel = 0

tape_address = /dev/rmt/0mn

tape_address_rew = /dev/rmt/0m

volume_archive = (SIDA01, SIDA02, SIDA03, SIDA04, SIDA05,SIDA06, SIDA07, SIDA08,
SIDA09, SIDA10,SIDA11, SIDA12, SIDA13, SIDA14, SIDA15,SIDA16, SIDA17, SIDA18, SIDA19,
SIDA20,SIDA21, SIDA22, SIDA23, SIDA24, SIDA25,SIDA26, SIDA27, SIDA28, SIDA29, SIDA30)

volume_backup = (SIDB01, SIDB02, SIDB03, SIDB04, SIDB05, SIDB06, SIDB07, SIDB08,
SIDB09, SIDB10,SIDB11,SIDB12, SIDB13, SIDB14, SIDB15,SIDB16,SIDB17, SIDB18, SIDB19,
SIDB20,SIDB21,SIDB22, SIDB23, SIDB24, SIDB25, SIDB26, SIDB27, SIDB28, SIDB29, SIDB30)


expir_period = 14

tape_use_count = 100


parallel_instances =(

                    C111:C:\Oracle\C11\12102@C111,

                    C112:c:\Oracle\C11\12102@C112,

                    C113:C:\Oracle\C11\12102@C113)
```

asm_ora_sid = +ASM1

asm_ora_home = C:\oracle\GRID\12102

asm_root_dir = ASM

The parameter 'asm_ora_sid' must be adjusted to '+ASM2','+ASM3',... to run archivelog backups on a cluster node other than node 1.

Database Backup:
```
backup_mode = full

restore_mode = all

backup_type = online

backup_dev_type = rman_disk

backup_root_dir = c:\oracle\C11\sapbackup

stage_root_dir = c:\oracle\C11\sapbackup

compress = no
```

```
compress_cmd = "compress -c $ > $"

uncompress_cmd = "uncompress -c $ > $"

compress_dir = c:\oracle\C11\sapreorg

archive_function = save_delete

archive_copy_dir = c:\oracle\C11\sapbackup

archive_stage_dir = c:\oracle\C11\sapbackup

tape_copy_cmd = cpio

disk_copy_cmd = copy

stage_copy_cmd = rcp

pipe_copy_cmd = rsh

cpio_flags = -ovB

cpio_in_flags = -iuvB

cpio_disk_flags = -pdcu

rman_channels = 4

rman_filesperset = 1

_rman_sess_stmt = ("alter session set optimizer_mode=RULE")

rman_send = (

"channel sbt_1 'NSR_ENV=(NSR_SERVER=xxbackapp02-
se,NSR_CLIENT=oracx3v1,NSR_DATA_VOLUME_POOL=IDES Oracle1 4
Weeks,NSR_GROUP=2200xxbackup02sn01ora1)'",

"channel sbt_2 'NSR_ENV=(NSR_SERVER=xxbackapp02-se,NSR_CLIENT=
oracx3v1,NSR_DATA_VOLUME_POOL=IDES Oracle2 4
Weeks,NSR_GROUP=2200xxbackup02sn01ora2)'",

"channel sbt_3 'NSR_ENV=(NSR_SERVER=xxbackapp02-se,NSR_CLIENT=
oracx3v1,NSR_DATA_VOLUME_POOL=IDES Oracle3 4
Weeks,NSR_GROUP=2200xxbackup02sn01ora3)'",

"channel sbt_4 'NSR_ENV=(NSR_SERVER=xxbackapp02-se,NSR_CLIENT=
oracx3v1,NSR_DATA_VOLUME_POOL=IDES Oracle4 4
Weeks,NSR_GROUP=2200xxbackup02sn01ora4)'"

)

rewind = "mt -f $ rewind"

rewind_offline = "mt -f $ offline"

tape_pos_cmd = "mt -f $ fsf $"

tape_size = 100G
```

```
exec_parallel = 0

tape_address = /dev/rmt/0mn

tape_address_rew = /dev/rmt/0m

volume_archive = (SIDA01, SIDA02, SIDA03, SIDA04, SIDA05, SIDA06, SIDA07, SIDA08,
SIDA09, SIDA10, SIDA11, SIDA12, SIDA13, SIDA14, SIDA15, SIDA16, SIDA17, SIDA18,
SIDA19, SIDA20, SIDA21, SIDA22, SIDA23, SIDA24, SIDA25, SIDA26, SIDA27, SIDA28,
SIDA29, SIDA30)

volume_backup = (SIDB01, SIDB02, SIDB03, SIDB04, SIDB05, SIDB06, SIDB07, SIDB08,
SIDB09, SIDB10, SIDB11, SIDB12, SIDB13, SIDB14, SIDB15, SIDB16, SIDB17, SIDB18,
SIDB19, SIDB20, SIDB21, SIDB22, SIDB23, SIDB24, SIDB25, SIDB26, SIDB27, SIDB28,
SIDB29, SIDB30)


expir_period = 14

tape_use_count = 100


parallel_instances = (

                 C111:c:\Oracle\C11\12102@C111,

                 C112:c:\Oracle\C11\12102@C112,

                 C113:c:\Oracle\C11\12102@C113)


asm_ora_sid = +ASM1

asm_ora_home = C:\oracle\GRID\12102

asm_root_dir = ASM
```

SAP Standalone Gateway

In order to be able to run automated database and archivelog backups from transaction 'DB13' install a SAP Standalone Gateway on each of the cluster nodes and extract the latest BR*Tools bundle.

Please refer to the following SAP notes for more information:

SAP Note #	
1025707	DBA Cockpit: Planning calendar and remote Oracle databases
853610	Configuring transaction DB13C for Oracle
1028624	Overview of DBA Cockpit for Oracle

Testing

Failures

One of the most important things before going into production with a new cluster environment is testing. Only by testing the behaviour of the cluster by simulating hardware and software failures ensures that it will react as expected when a real failure occurs. The following two sections describe some essential tests. It is highly recommended not to go into production if those tests do not show the expected results. The list is not complete and customers may think about additional tests that have to be done for their specific cluster environment and HA requirements.

Hardware Failures

Power Loss or Complete Failure of Cluster Node.

Action

One of the most likely failures that may occur is a crash of a cluster node. Simulate a crash by issuing a "hard power-off" for one of the cluster nodes.

Expected Behaviour

One of the remaining nodes should do a crash recovery for the crashed node. Depending on the service definitions and their failover policy the services may be activated on one of the remaining nodes. The whole failover process may take some minutes. To verify whether the failover process is finished run 'crsctl status resource -t'. Check that one of the remaining nodes finished the crash recovery by viewing the alert.logs. Finally the VIP of the failed cluster node and all services should be available as defined.

Action

Power on the shutdown cluster node.

Expected Behaviour

The VIP of the started cluster node that was serviced by one of the remaining cluster nodes should fall back to the rebooted cluster node. All defined Oracle instances (if more than one database instance was running on the node before) should start automatically. All services that fell over to other cluster nodes should stay there. They have to be moved manually to their corresponding nodes using 'srvctl relocate service...'.

Oracle Interconnect Failures

Cluster node has lost its connection to the switch.

Action

Disconnect one cluster node from the switch by unplugging all network cables of the Oracle Interconnect. If you have NIC teaming active (redundant interconnect) also check what happens when you unplug only one network cable.

Expected Behaviour

In a two node RAC cluster one of both nodes should crash. In a three or more node RAC cluster the node where you did unplug the network cables should crash. In both situations the VIP should move to one of the remaining nodes. One of the remaining instances should do the crash recovery for the crashed node. Services should move to the remaining nodes as defined.

(Multiple) Switch Failures with all nodes losing their connection.

Action

Turn off or disable the network switch of the Oracle Interconnect.

Expected Behaviour

Regardless of the number of cluster nodes all nodes except one should crash. The remaining node should take over the VIP´s of the crashed nodes, perform a crash recovery and provide the database services as defined.

Storage Failures

One cluster node loses its storage connection.

Action

Unplug one of the cluster nodes from the storage. If you have redundant paths to the storage also check what happens when you unplug only one (fiberchannel-) cable.

Expected Behaviour

The node should crash without automatic reboot. Automatic reboot may be a configuration option in the BIOS of the node.

All cluster nodes lose their storage connection.

Action

Unplug all cluster nodes from the storage or turn off all switches to the storage.

Expected Behaviour

All cluster nodes should crash. After reconnecting the nodes to the storage and rebooting all the nodes the database should be available as expected.

Software Failures

Shutdown abort.

Action

Shutdown an Oracle instance with "shutdown abort".

Expected Behaviour

The cluster node should stay up and running. One of the remaining Oracle instances should do the crash recovery. All services of the stopped instance should move to the previously defined remaining instances. 'crsctl status resource -t' utility should display "OFFLINE" as the target state for the stopped instance.

Kill oracle.exe

Action

Kill oracle.exe with the Windows Task Manager.

Expected Behaviour

The cluster node should stay up and running. One of the remaining Oracle instances should do the crash recovery. All services of the killed instance should move to the previously defined remaining instances. CRS should start the killed instance automatically. All services that moved to other instances should stay there. They have to be relocated manually.

SAP Application Workload

The most important tests to run are workload tests where customers should try to run their SAP applications under simulated business conditions. Testing the performance of core transactions and batch jobs will help to identify and resolve bottlenecks and configuration issues before going into production with your new cluster environment. These tests can be combined with some of the failure test scenarios described earlier.

Check if local update dispatching (described in the next chapter) is active and working as expected with "dpmon.exe" or SAP Management Console. Both tools show the number of times an update process was active.

Check if database service relocation works as expected and every SAP application server can connect to its service.

Additional Performance Recommendations

Using ASSM tablespaces is mandatory. If you still have Non-ASSM tablespaces you must migrate your tables to ASSM tablespaces first.

Use local update dispatching for SAP update-processes

Local Update Dispatching

The SAP application performs updates on the database asynchronously through 'update-processes'. A dialog work process inserts an update job into tables VBHDR, VBDATA, VBMOD and posts the dispatcher process on the central instance to select an update-process to actually perform the update job. This update-process may be connected to another RAC instance, similarly to the dialog-process. If this is the case, the update-process would read the data just written by the dialog-process. To satisfy this read request, all required database blocks have to be shipped over the interconnect from the instance where the data was inserted to the instance where the data has to be read. This approach produces unnecessary additional data load and can be avoided.

In order to avoid these 'non-local updates', it is recommended that you:

• 	have a number of update-processes on each SAP instance

• 	turn off update-dispatching and use the local update-processes

To turn off update dispatching, set rdisp/vb_dispatching to 0 as described in SAP note #62203. Also, set rdisp/vb_name to the name of the local instance (e.g. rdisp/vb_name = app1_C11_00).

Integrated Cloud Applications & Platform Services